

Policy-Based File Tiering

Reduce Costs and Reclaim Capacity from Your Windows File Shares by Offloading Less Relevant Files to On-Premises Swarm Object Storage and the Cloud

BENEFITS

- Free up precious capacity on your primary storage for more critical data
- Defer purchase of additional premium storage
- Shorten backup cycles and reduce load on your existing storage
- Eliminate time wasted performing manual file migrations
- See the space and cost savings potential upfront with “what-if” analysis

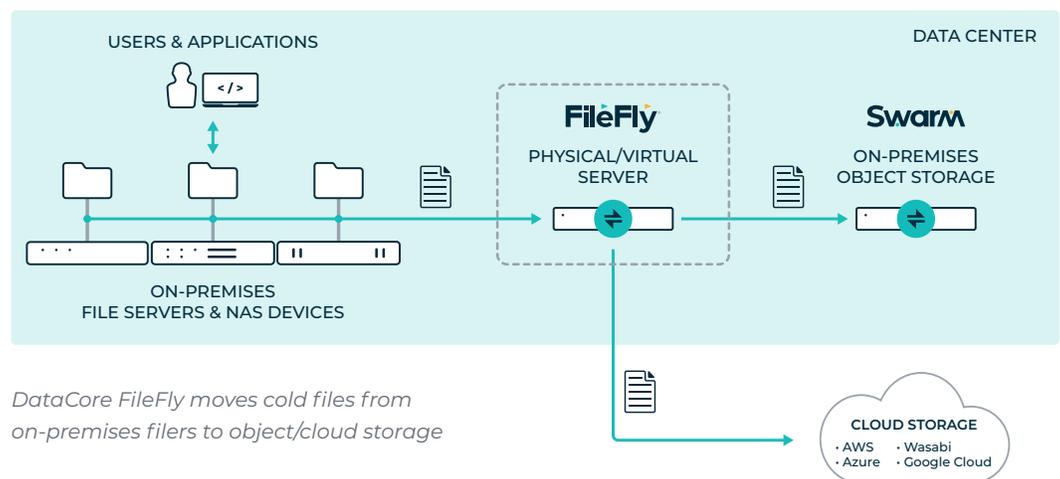
We instinctively know that much of the storage space in our servers and NAS is wasted on files that no one really cares much about. Yet deleting them is not an option. They may be necessary to comply with data retention regulations and to recall records of past activities.

FileFly software periodically migrates the contents of files you deem less relevant to cheaper secondary storage tiers, constantly making room for important new data. The migrated files remain accessible from the same folders as before. In the rare case that the file is requested, its data is transparently recalled to primary storage.

Simple rules and policies determine which files are migrated and when the data transfers should take place. Characteristics such as age, owner, file type and size may be used to include or exclude certain files. Scans for files meeting the criteria take place on schedule or on-demand. Policy changes can be simulated to see their effect on space savings and cost reduction before performing the actual tiering operations.

HOW FILE-BASED TIERING WORKS

FileFly regularly scans selected Windows shares and subdirectories, looking through file metadata for attributes that match the selection criteria. A migration task copies the file’s contents to secondary storage, then marks the original file with a stub pointing to the new data location. The original file is truncated to zero physical size and tagged as offline, freeing up its primary storage space. The logical size and other original attributes, including access controls and ownership, remain for the correct operation of applications and users. The offline flag serves two purposes. It alerts users that file retrieval may take a little longer. And it signals backup products to skip those files. Best practices call for files to be backed up before migration.



When a previously migrated file is opened, its data will be recalled back to the original source filer. If no changes are made, the copy in secondary storage is unaffected. This eliminates the need for a data transfer during a “quick re-migrate” operation. For files that have been modified or deleted, periodic “scrub” tasks delete outdated copies to reclaim capacity in secondary storage.

Migrate mp4 video files older than 3 years owned by Juergen and Marty who are no longer working here

Negate

Filename patterns
*.mp4

Minimum size: 0 bytes | Maximum size: Unlimited bytes

Date matching
Modified Older than 3 Years

Owner patterns
aglab.local\juergen, aglab.local\marty
Leave empty to match all users
eg. SomeDomain\joe
eg. SomeDomain* (wildcard matching)
eg. /SomeDomain/[a-k].* (regular expression matching)

Attribute state matching

The screenshot shows a dashboard with several panels: 'Notices' (warning about servers), 'Operations' (a diamond-shaped diagram with 'Migrate', 'Quick Re-migrate', 'Change Tier', 'Pre-migrate', and 'Scrub' nodes), 'Servers' (a gauge showing 1 warning and 0 errors), 'Quota' (a gauge showing 281.4 MB / 10 TiB (00% remaining)), 'Primary' and 'Secondary' storage usage bar charts, 'Processed' (a line graph showing thousands per hour), and 'Operations' (a line graph showing ops per day). A 'Recent Tasks' panel at the bottom shows a task 'gather stats for marty winserv-marketing'.

Define file tiering rules and track progress for multiple Windows file servers from a central console

The analysis, charting, migration, and policy-based file tiering features of FileFly result in several economic benefits:



DEFER ADDITIONAL HARDWARE PURCHASES

of expensive primary file storage



REDUCE PRICE/TB

for retaining less relevant data on lower cost storage



FREE UP VALUABLE ADMINISTRATION TIME

by automating file migrations according to custom policies



SAVE ON BACKUP

resources and reduce backup windows

LICENSING

Benefit from simple, transparent, and flexible licensing according to how many Terabytes (TBs) of data are migrated FROM source file servers and NAS devices TO object storage (on-premises or cloud).

Pricing includes 24x7 Premier Support and software updates, and price per TB goes down as consumption grows (volume discounts).

SPECIFICATIONS

PRIMARY FILE SYSTEMS (SOURCES)

- Windows File Servers: NTFS and SMB Shares
- SMB Shares on NetApp and Dell EMC Isilon NAS

SECONDARY STORAGE TIERS (DESTINATIONS)

- On-premises: DataCore Swarm object storage
- Cloud: S3-compatible storage such as AWS Amazon S3, Microsoft Azure, Google Cloud, Wasabi Hot Cloud and other qualified S3-compatible providers

FILEFLY COMPONENTS

In its simplest form, FileFly consists of three software components:

- The web-based Admin Portal for centralized control of configurations, task scheduling, monitoring and reporting. It lies outside the data path.
- A lightweight Migration Agent installed on each Windows Server hosting NTFS file systems.
- A Gateway Agent installed on a separate Windows Server instance (VM or bare metal) responsible for transferring data to and from secondary storage. Redundant gateway instances may be configured for high availability.

To connect to SMB Shares on NetApp and Dell EMC Isilon NAS:

- NetApp NAS systems require an Fpolicy Server (two for HA)
- Isilon NAS systems require a LinkConnect Server instance (two for HA). Each Windows Client also needs a lightweight LinkConnect Driver.

| Minimum Requirements | For Windows Server Components |
|---------------------------------------|---|
| Operating Systems | Windows Server 2019, 2016, 2012 R2, 2012 |
| CPU (Intel/AMD x86 processors) | 2 vCPU cores |
| Memory | 4 GB |
| Disk space | 2 GB for log files (Additional space for LinkConnect Cache) |
| Network | 1 Gbit/s Ethernet |



Discover the Ultimate Flexibility of DataCore Software

DataCore Software delivers the industry's most flexible, intelligent, and powerful software-defined storage solutions for block, file, and object storage, helping more than 10,000 customers worldwide modernize how they store, protect, and access data. With a comprehensive product suite, intellectual property portfolio, and unrivaled experience in storage virtualization and advanced data services, DataCore is The Authority on Software-Defined Storage. www.datacore.com

GET STARTED